

## **Transparent Memory Address Remapping**

### **CROSS REFERENCE TO RELATED APPLICATIONS**

**[0001]** This application is a continuation and claims priority of pending U.S Patent  
5 Application 10/124,783, filed 17 April 2002.

### **BACKGROUND OF THE INVENTION**

#### **FIELD OF THE INVENTION**

**[0002]** This invention relates to the field of memory management in computers, in  
10 particular in the context of address mapping in order to improve I/O speed.

#### **DESCRIPTION OF THE RELATED ART**

**[0003]** Many computer systems depend for their speed and efficiency on the ability  
to rapidly transfer data between devices and system memory. In many cases, however,  
15 addressing conventions and restrictions make it necessary to perform intermediate  
copies of data to be transferred before the final transfer can actually take place. Such  
copying can severely slow down the transfer rate.

**[0004]** One widely used method for increasing the input/output ("I/O" -- either or  
both) speed between certain devices (or other processes) and memory is known as  
20 "direct memory access" (DMA). DMA is a capability provided by some computer bus  
architectures that allows data to be sent directly from an attached device (such as a disk  
drive) to system memory, without intermediate action by the processor. In order to  
implement DMA, a portion of system memory is usually designated as an area to be  
used specifically for DMA operations. Obviously, time is lost whenever a block of data  
25 (such as a "page" that is not already in the designated memory portion) must be copied  
to or from the designated memory portion to perform a DMA transfer.

**[0005]** As a concrete example, modern Intel x86 processors support a physical  
address extension (PAE) mode that allows the hardware to address up to 64 GB of  
memory using 36-bit addresses. Unfortunately, many devices that directly access  
30 memory to perform I/O operations can address only a subset of this memory. For  
example, network interface cards with the common 32-bit PCI (Peripheral Component

Interconnect) interface can address memory residing in only the lowest 4 GB of memory, even on systems that support up to 64 GB of memory. Other 32-bit PCI devices can access memory above 4GB using a technique known as DAC (Dual Address Cycle), but this technique requires two address transfers -- one for the low 32 bits and another for the high 32 bits.

**[0006]** One known way to support output to "high" memory (that is, memory above 4 GB) is to copy the data from high memory to a temporary buffer in "low" memory for the DMA operation. For input operations, a portion of low memory in the temporary buffer is allocated for storage of the input data, which can then be copied to high memory. This technique is employed, for example, by the Linux 2.4 kernel, which uses the term "bounce buffer" to describe the temporary buffering and copying process. Unfortunately, copying can impose significant overhead, which results in turn in increased latency, reduced throughput, and/or increased CPU load when performing I/O.

**[0007]** Another known technique is the remapping of memory regions (in particular, pages) as described in U.S. Patent No. 6,075,938, Bugnion, et al., "Virtual Machine Monitors for Scalable Multiprocessors," issued 13 June 2000 ("Bugnion '938"). The basic idea of this system, which operates in the context of a NUMA (non-uniform memory access) multi-processor, is that memory pages associated with hardware memory modules that are farther away (defined in terms of access latency) are migrated or replicated by making copies in hardware memory modules closer to a process that is accessing them. The process page mappings are modified transparently to use the local page copy instead of the original remote page. In other words, the Bugnion '938 system attempts to improve access speed by improving memory locality. The problem when it come to I/O, in particular in the context of DMA, is, however, not that of whether a certain memory space is sufficiently local, but rather, often, whether it can be accessed at all.

**[0008]** Still other existing systems enable I/O to "high" memory by including special hardware components that provide support for memory remapping. For example, a separate I/O memory management unit (I/O MMU) may be included for I/O operations. The obvious disadvantage of this solution is its requirement for the extra hardware.

**[0009]** A related problem is the dynamic management of the "low" memory, which may be a scarce resource that needs to be allocated among various competing uses.

In other words, if several devices or processes must compete for use of a common

5 memory region (here, "low") designated for high-speed I/O (such as DMA), then some mechanism must be provided to efficiently allocate its use. Such memory management is typically carried out by a component of the operating system.

**[0010]** What is needed is therefore a system that eliminates or at least reduces the need for copying in I/O operations to or from at least one limited memory space,

10 especially in high-speed I/O contexts such as DMA. The system should preferably be usable not only in a conventional computer system, in particular, in its operating system, but also in computer systems that include at least one virtualized computer. Moreover, the system should preferably also be able to manage the limited memory space dynamically, and it should not require specific hardware support. This invention

15 provides such a system and method of operation whose various aspects meet these different goals.

## SUMMARY OF THE INVENTION

**[0011]** The invention provides a method for performing an input/output (I/O)

20 operation in a computer between an I/O-initiating subsystem and a device through a memory, in which the memory is arranged into portions that are separately addressable using first identifiers that are represented using a first number of bits; for the I/O

operation, the device accesses a first space of the memory; and the subsystem addresses I/O requests to a second space of the memory using second identifiers that

25 are represented using a second number of bits. The second identifiers are initially mapped to respective first identifiers that identify portions of the memory in the second memory space. For any I/O request that meets a remapping criterion, the corresponding second identifier is remapped to one of the first identifiers that identifies a portion of the memory in the first space of the memory. The second space is different  
30 from the first space and the second number of bits is greater than the first number of bits.

**[0012]** Each first identifier may be generated to have a subset of bits identical to corresponding bits of the second identifier during remapping.

**[0013]** According to one optional aspect of the invention, a new copy of the data set in the buffer upon each instance of the I/O request is created for any I/O request that fails to meet the remapping criterion.

**[0014]** According to another optional aspect of the invention, a new copy of the data set in the buffer upon each instance of the I/O request is created for any I/O request that fails to meet the remapping criterion.

**[0015]** According to still another optional aspect of the invention, for each second identifier that is currently mapped into the first space of the memory and that meets a remapping condition, the second identifier is again remapped into the second space of the memory. The portion of the memory in the first space to which the second identifier had previously been remapped may then be freed for reallocation.

#### BRIEF DESCRIPTION OF THE DRAWINGS

**[0016]** Figure 1 illustrates a generalized embodiment of the invention in the case where memory addresses issued by an I/O-initiating subsystem undergo a single mapping in order to address the actual system memory for the purpose of I/O.

**[0017]** Figure 2 illustrates a generalized embodiment of the invention in the case where memory addresses issued by an I/O-initiating subsystem undergo at least two mappings before they are used to address the actual system memory.

**[0018]** Figure 3 is a block diagram of the preferred embodiment of the invention, in which I/O via system memory is initiated by an I/O-initiating subsystem running on a guest operating system within a virtual machine, which in turn is running on an underlying host platform.

## DETAILED DESCRIPTION

**[0019]** Figure 1 shows a generalized embodiment of the invention, and also serves to illustrate the concepts used in other embodiments of the invention as well. In Figure 1, one or more subsystems 30, ..., 31, implemented in either software or hardware or a combination of both, initiates an I/O request to a device 400, which is able to access only a portion of the total memory space of a memory 112. Of course, as in other computer systems, the system according to the invention includes a hardware platform with one or more processors and supporting components, as well as system software such as an operating system (OS), etc.; these known system components (both hardware and software) are not shown in Figures 1 and 2 for the sake of simplicity but can be assumed to be present.

**[0020]** In most applications of the invention, it is anticipated that the device 400 will be one or more physical devices. The invention may also be used, however, even where the "device" is a software construct, that is, a virtualized device.

**[0021]** In the following description of the invention, the portion of the memory space that the device 400 can access is referred to as the "low" memory Mem\_L, with the remaining memory space referred to as "high" memory Mem\_H, only for the sake of ease of understanding and illustration. Note that this nomenclature also corresponds to the example mentioned above, namely, the case in which the device is a network interface card that can address memory residing in only the lowest 4 GB of memory, even on systems that support much larger memory address spaces.

**[0022]** In general, the invention is applicable whenever the addressable space of the device is not identical to the addressable space of the subsystem that is requesting I/O. This may be the case for any of several reasons. For example, the subsystem may be restricted to accessing some portion of high memory either by its own design, or in order to maintain isolation from other subsystems. Another example is where the device with which the subsystem believes it is interacting is actually an emulation that is addressed in high memory, with an intermediate subsystem acting as an interface to the physical device 400, which addresses only low memory; this example is explained in greater detail below.

**[0023]** The memory space addressable by the device need not be ordered numerically "lower" than the memory space that it cannot address, and it need not be contiguous, although this will be the most common arrangement where the operation to be supported is DMA. The term "low memory" is therefore to be understood simply as the portion of the memory that the device 400 is able to address; the term "high memory" is the portion that the device cannot or for some other reason does not address, but that the subsystem requesting a current I/O operation can address (directly or indirectly).

**[0024]** As is well known, in most modern computer architectures, system memory is typically divided into individually addressable units or blocks commonly known as "pages," each of which in turn contains many separately addressable data words, which in turn will usually comprise several bytes. In Intel x86 systems, for example, each page comprises 4096 bytes. Pages are identified by addresses commonly referred to as "page numbers." The invention does not presuppose any particular page size.

**[0025]** In broadest terms, the invention provides a mechanism that takes a request for I/O of data residing on one or more pages, and under certain conditions remaps the request from high memory to low memory. Thanks to the remapping procedure according to the invention, there is no need to create a copy of the page(s) for each I/O operation, but rather only a single copy need be created as long as remapping is in effect; no copying may be required at all for some read operations.

**[0026]** Note that the actual data involved in a given I/O operation need not occupy an entire page. In other words, the I/O granularity may be smaller than the granularity of memory allocation. Three of the many cases in which this is true are DMA, which may involve data blocks as small as single data words or even bytes; packets in a packet-based network transfer, which commonly require fewer than 1600 bytes each out of a standard 4096-byte page; and disk transfers, which typically take place in 512-byte sectors. Even in these cases, in which the I/O data occupies only a subset of a page, a page is usually the smallest unit of memory that can be remapped.

**[0027]** As Figure 1 shows, the invention includes a subsystem, shown as the "manager" 605, which acts as an interface or intermediary between the subsystem that is requesting I/O on the one hand, and the memory 112 and the device 400 on the other

hand. It is assumed that the manager 605 is able to detect which pages the subsystem 30 is designating for a current I/O operation with the device 400. In embodiments of the invention such as the one shown in Figure 1, the manager 605 will typically be a component of the operating system; an alternative configuration is described below.

5

## MEMORY MAPPING AND ADDRESS TERMINOLOGY

**[0028]** The most straightforward way for all components in a computer to uniquely identify a memory page would be for them all simply to use a common set of page numbers. This is almost never done, however, for many well-known reasons. Instead, user-level software normally refers to memory pages using one set of identifiers, which is then ultimately mapped to the set actually used by the underlying hardware memory.

10

**[0029]** When a subsystem 30 requests access to the memory 112, for example, the request is issued usually with a "virtual address," since the memory space that the subsystem addresses is a construct adopted to allow for much greater generality and flexibility. The request must, however, ultimately be mapped to an address that is issued to the actual hardware memory. This mapping, or translation, is typically specified by the operating system (OS). The OS thus converts the "virtual" page number (VPN) of the request into a "physical" page number (PPN) that can be applied directly to the hardware.

15

**[0030]** For example, when writing a given word to a virtual address in memory, the processor breaks the virtual address into a page number (higher-order address bits) plus an offset into that page (lower-order address bits). The virtual page number (VPN) is then translated using mappings established by the OS into a physical page number (PPN) based on a page table entry (PTE) for that VPN in the page table associated with the currently active address space. The actual translation may be accomplished simply by replacing the VPN (the higher order bits of the virtual address) with its PPN mapping, leaving the lower order offset bits the same.

20

**[0031]** Normally this mapping is performed within a hardware memory management unit (MMU) 116 and is obtained quickly by looking it up in a hardware structure known as a translation lookaside buffer (TLB); if not, a "TLB miss" occurs, and the page tables in memory are consulted to update the TLB before proceeding. The operating system

25

30

thus specifies the mapping, but the hardware MMU 116 usually actually performs the conversion of one type of page number to the other. Below, for the sake of simplicity, when it is stated that a software module “maps” pages numbers, the existence and operation of a hardware device such as the MMU 116 may be assumed.

5   **[0032]**   The concepts of VPNs and PPNs, as well as the way in which the different page numbering schemes are implemented and used, are described in many standard texts, such as "Computer Organization and Design: The Hardware/ Software Interface," by David A. Patterson and John L. Hennessy, Morgan Kaufmann Publishers, Inc., San Francisco, California, 1994, pp. 579-603 (chapter 7.4 "Virtual Memory"). In the  
10   generalized embodiment of the invention shown in Figure 1, the manager 605 includes a mapping module 610 that includes a map 612 of VPNs to PPNs. For each VPN issued by the subsystem 30, the manager can therefore determine which corresponding PPN is to be used to address the physical memory 112.

**[0033]**   Figure 2 illustrates a system in which memory is addressed using a second  
15   level of indirection, that is, where the VPN issued by a subsystem 530 is remapped twice in order to determine which page of the hardware memory is intended. In this embodiment, the subsystem 530 is included within a guest system 500, which includes its own guest operating system 520 or equivalent system software, as well as a guest memory 512, which may in actual implementations be a designated subset of the  
20   memory space of the hardware memory 112. The guest assumes that the device 540 with which the subsystem 530 is attempting an I/O operation is the device that will actually carry out the intended I/O operation, and it assumes that the device 540 addresses the guest memory 512.

**[0034]**   In other words, at least with respect to I/O between the subsystem 530 and  
25   the device 540, the guest 500 acts as a normal computer system. In fact, however, the device 540 is either a software construct, such as an emulator, or for some other reason not the actual hardware device 400 that is to carry out the requested I/O operation. When the subsystem requests I/O of data residing on a page, it therefore designates a VPN as usual. A mapping module 510 within the guest OS then translates this VPN  
30   into a corresponding PPN in the conventional manner and then uses this PPN as usual to address the memory 512. The guest OS therefore “believes” that it is directly



addressing the actual hardware memory via which the I/O is to occur, but in fact it is not. Of course, since the device 540 is not a “real” I/O device, actual transfer must ultimately be arranged with the hardware device 400.

**[0035]** In the embodiment shown in Figure 2, it is assumed that all I/O requests of the guest system 500, at least those involving the device 400, are detected and handled by the manager 605. The manager 605 therefore intercepts the PPN issued by the mapping module 512 (or the original VPN) and maps it to the actual page number used by the hardware memory. In this multiple-mapped embodiment, the page numbers used to address the hardware memory 112 are referred to as “machine page numbers” (MPNs). The memory mapping module 612 in the manager 605 is therefore shown in Figure 2 as mapping PPNs to MPNs.

**[0036]** Figure 2 therefore illustrates a configuration in which the guest system 500 may be entirely a software construct, or at least a system whose actual execution is handled in whole or in part by an underlying software and hardware platform. In other words, at least some of the hardware structures that the guest OS assumes are handling the actual I/O operation are in fact virtualized. In the preferred embodiment of the invention described below, the guest system is virtual machine, with all essential hardware structures implemented in software. Any single-mapping embodiment of the invention, which has the general structure shown in Figure 1, is therefore referred to here as “non-virtualized” embodiment or configuration, whereas embodiments having the general structure shown in Figure 2 are referred to as “virtualized” embodiments or configurations. Note that from the perspective of the manager 605, the guest system 500 can be considered to be a subsystem that is issuing the I/O request, even though it may carry out internal address mappings of its own.

**[0037]** Because of the additional degree of addressing indirection introduced by virtualization, the concept of a physical page number (PPN) differs in non-virtualized and virtualized embodiments of the invention. In particular, in virtualized embodiments of the invention, the definition of a PPN deviates from the “standard” definition. As used in herein, the definitions of VPN, PPN and MPN are as follows:

#### Non-Virtualized Configuration:

**[0038]** VPN: A virtual page number associated with a process.

**[0039]** PPN: A physical page number that refers to the actual hardware memory.

The operating system, in particular, the mapping module 610, specifies mappings from

5 VPNs to PPNs, and the hardware MMU 116 then performs the actual translation of VPNs to PPNs using these mappings.

**[0040]** MPN: A machine page number, identical to PPN.

#### Virtualized Configuration:

10 **[0041]** VPN: A virtual page number associated with a subsystem running in or on a guest OS.

**[0042]** PPN: A physical page number that refers to a virtualized physical memory space associated with the guest. As is mentioned above, the guest operates as though this PPN refers to actual hardware memory, although it is actually a software construct  
15 maintained by the guest software layer. The guest OS specifies mappings from VPNs to PPNs.

**[0043]** MPN: A machine page number that refers to actual hardware memory 112. The intermediate software layer (for example, a virtual machine monitor VMM acting as the manager 605) specifies mappings from each VM's PPNs to MPNs. This adds an  
20 extra level of indirection, with two address translations (mappings) instead of one: a VPN is translated to a PPN using the guest OS mappings, and then this PPN is mapped to an MPN by the manager. In order to eliminate one mapping operation while still maintaining the extra degree of addressing indirection, the manager may instead (or, if needed, in addition) maintain a separate page table from VPNs to MPNs, so that the  
25 hardware MMU 116 can translate VPNs directly to MPNs, and remap them as described below.

**[0044]** Note that the page-mapping operation within the manager 605 may be kept transparent to the guest 500 that is requesting I/O. Consequently, as long as the corresponding I/O operation is carried out, the guest system is unaware of and  
30 unaffected by the mapping used to address the hardware memory.

## PREFERRED VIRTUALIZED EMBODIMENT

[0045] Figure 3 illustrates the main components of the preferred, virtualized embodiment of the invention, in which the guest is a software construct known as a "virtual computer" or "virtual machine" (VM) 500 (a special case of the guest 500 shown in Figure 2). Reference numerals in Figure 3 that are the same as in Figure 2 therefore refer to corresponding components and features. As in conventional computer systems, the virtualized embodiment includes both system hardware 100 and system software 200. The system hardware 100 includes one or more central processors CPU(s) 110, which may be a single processor, or two or more cooperating processors in a known multiprocessor arrangement. As in most computers, one or more disks 114 are usually included in addition to the system memory 112. The system hardware usually also includes, or is connected to, conventional registers, interrupt-handling circuitry, a clock, etc., as well as the memory management unit MMU 116.

[0046] The device 400 that is to directly access the memory 112 in accordance with an I/O request from the VM is either part of the system hardware 100, or is connected to the hardware 100 as a peripheral. Just two of the many possible examples of devices that may be configured for DMA according to the invention are the disk 114 itself (which is an example of a device that is also part of the system hardware) and a network interface card.

[0047] The system software 200 either is or at least includes an operating system OS 220, which will include drivers 222 as needed for controlling and communicating with various devices, usually including the disk 114. Conventional applications 300, if included, may be installed to run on the hardware 100 via the system software 200.

[0048] As is well known in the art, a virtual machine (VM) is a software abstraction -- a "virtualization" -- of an actual physical computer system. As such, each VM will typically include virtualized ("guest") system hardware 501 and guest system software 502, which are software analogs of the physical hardware and software layers 100, 200. Note that although the hardware "layer" 501 will be a software abstraction of physical components, the VM's system software 502 may be the same as would be loaded into a "real" computer. The modifier "guest" is used here to indicate that the various VM software components, from the perspective of a user, are independent, but that actual

execution is carried out on the underlying "host" hardware and software platform 100, 200. The guest system hardware 501 includes one or more virtual CPUs 510 (VCPU), virtual system memory 512 (VMEM), a virtual disk 514 (VDISK), and at least one virtual device 540 (VDEVICE), all of which are implemented in software to emulate the  
5 corresponding components of an actual computer. Of particular relevance here is that, from the perspective of the VM, the virtual device 540 is a software analog of the hardware device 400. Thus, I/O to the virtual device will actually be carried out by I/O to the hardware device 400, but in a manner that is transparent to the VM.

**[0049]** The guest system software 502 includes a guest operating system 520, which  
10 may, but need not, simply be a copy of a conventional, commodity OS, as well as drivers 522 (DRVS) as needed, for example, to control the virtual device 540. The guest OS also includes the page map 510.

**[0050]** Of course, most computers are intended to run various applications, and a VM is usually no exception. Consequently, by way of example, Figure 3 illustrates one  
15 or more applications 503 installed to run on the guest OS 520; any number of applications, including none at all, may be loaded for running on the guest OS, limited only by the requirements of the VM. The subsystem that issues the I/O request and the corresponding VPN(s) may be any of the applications 503, some subsystem within the guest OS 520, or possibly some other virtual device than device 540.

**[0051]** If the VM is properly designed, then the applications (or the user of the  
20 applications) will not "know" that they are not running directly on "real" hardware. Of course, all of the applications and the components of the VM are instructions and data stored in memory, just as any other software. The concept, design and operation of virtual machines are well known in the field of computer science. Figure 3 illustrates a  
25 single VM 500 merely for the sake of simplicity; in many installations, there will be more than one VM installed to run on the common hardware platform; all will have essentially the same general structure, although the individual components need not be identical.

**[0052]** Some interface is usually required between the VM 500 and the underlying  
"host" hardware 100, which is responsible for actually executing VM-related instructions  
30 and transferring data to and from the actual, physical memory 112. One advantageous interface between the VM and the underlying host system is often referred to as a virtual

machine monitor (VMM). Virtual machine monitors have a long history, dating back to mainframe computer systems in the 1960s. See, for example, Robert P. Goldberg, "Survey of Virtual Machine Research," IEEE Computer, June 1974, p. 54-45. A VMM is usually a relatively thin layer of software that runs directly on top of a host, such as the system software 200, or directly on the hardware, and virtualizes the resources of the (or some) hardware platform. The VMM will typically include at least one device emulator 640, which may also form the implementation of the virtual device 540. The interface exported to the respective VM is usually such that the guest OS 520 cannot determine the presence of the VMM. The VMM also usually tracks and either forwards (to the host OS 220) or itself schedules and handles all requests by its VM for machine resources as well as various faults and interrupts. The general features of VMMs are known in the art and are therefore not discussed in further detail here.

**[0053]** In Figure 3, a single VMM 600 is shown acting as the interface for the single VM 500. It would also be possible to include the VMM as part of its respective VM, that is, in each virtual system. Although the VMM is usually completely transparent to the VM, the VM and VMM may be viewed as a single module that virtualizes a computer system. The VM and VMM are shown as separate software entities in the figures for the sake of clarity. Moreover, it would also be possible to use a single VMM to act as the interface for more than one VM, although it will in many cases be more difficult to switch between the different contexts of the various VMs (for example, if different VMs use different guest operating systems) than it is simply to include a separate VMM for each VM. This invention described below works with all such VM/VMM configurations.

**[0054]** The important point is simply that some well-defined interface should be provided between each installed VM 500 and the underlying system hardware 100 and that this interface should contain a manager 605 that is structured and functions as the manager 605 in Figure 2. Consequently, instead of a complete VMM, in order to implement the invention, the manager 605 (in Figure 2, manager 605) may be included in the host OS 220 or in some other system-level software layer. One advantage of including the manager 605 in the VMM is that this will not require modification of the host OS 220.

**[0055]** In some configurations, the VMM 600 runs as a software layer between the host system software 200 and the VM 500. In other configurations, such as the one illustrated in Figure 3, the VMM runs directly on the hardware platform 100 at the same system level as the host OS. In such case, the VMM may use the host OS to perform certain functions, including I/O, by calling (usually through a host API – application program interface) the host drivers 222. In this situation, it is still possible to view the VMM as an additional software layer inserted between the hardware 100 and the guest OS 520. Furthermore, it may in some cases be beneficial to deploy VMMs on top of a thin software layer, a "kernel," constructed specifically for this purpose.

**[0056]** In yet other implementations, the kernel takes the place of and performs the conventional functions of the host OS. Compared with a system in which VMMs run directly on the hardware platform, use of a kernel offers greater modularity and facilitates provision of services that extend across multiple virtual machines (for example, resource management). Compared with the hosted deployment, a kernel may offer greater performance because it can be co-developed with the VMM and be optimized for the characteristics of a workload consisting of VMMs.

**[0057]** As used herein, the "host" OS therefore means either the native OS 220 of the underlying physical computer, or whatever system-level software handles actual I/O operations, takes faults and interrupts, etc. for the VM. The invention may be used in all the different configurations described above.

**[0058]** In addition to controlling the instruction stream executed by software in virtual machines, the VMM also controls other resources in order to ensure that the virtual machines remain encapsulated and do not interfere with other software on the system. First and foremost, this applies to I/O devices, but also to interrupt vectors, which generally must be directed into the VMM (the VMM will conditionally forward interrupts to the VM). Furthermore, the memory management (MMU) functionality normally remains under control of the VMM in order to prevent the VM from accessing memory allocated to other software on the computer, including other VMs. In short, the entire state of the VM is not only observable by the VMM, but the entire operation of the VM is under the control of the VMM.

## REMAPPING

**[0059]** As Figures 1 and 2 show, the mapping module 610 in the manager 605 (or in the VMM 600 -- Figure 3) according to the invention includes a page activity module 614 and a remapping module 616, which together implement the main novel features of the invention. The manager preferably also includes a page copy buffer 620, which may be implemented as an explicitly-designated structure in low memory, or simply by allocating memory out of a general, limited "low memory pool" to perform remapping.

**[0060]** This buffer 620 is used as in conventional systems: as needed, I/O to "high" memory is supported by the manager temporarily copying data into the low-memory buffer 620. The buffer 620 is therefore equivalent to the "bounce buffer" described above. Unlike in the prior art, however, the buffer as used in this invention is not always needed for "high-memory" page I/O. The size of the buffer 620 may be fixed at a size determined using normal design methods, or it may be allowed to grow and shrink as a function of current need and low-memory availability. Recall that, for output operations, data will be copied into the buffer page(s) from the subsystem requesting the operation to the device, whereas, for input operations, the page(s) will be allocated for storage of data received from the device, after which the pages will be available for access by the requesting subsystem. Note that at least a partial page copy may be needed for input (read) operations in that case where the read involves only part of a page: A copy should then be made of at least the rest of the page, that is, the part not involved in the read operation.

**[0061]** The main idea of the invention is that a page that is frequently involved in I/O operations, and that would normally need to be copied through the buffer 620 upon each I/O operation involving that page, is instead transparently remapped into low memory. In other words, the VPN→PPN (non-virtualized embodiment) or the PPN→MPN (virtualized embodiment) mapping for highly active or "hot" pages is changed within the manager 605 so as to map directly into low memory.

**[0062]** In the virtualized embodiment of the invention, if a separate VPN→MPN map is maintained in the manager, then remapping may be done by changing this mapping instead, or in addition, to the PPN→MPN mapping. Because this mapping/remapping takes place at a functional level between that of the guest and host systems, whichever

set of page numbers (VPN or PPN) are used in the remapping process, they can be considered to be "intermediate" identifiers of the pages involved. Remapping according to the invention is discussed below in the context of changing PPN→MPN mappings simply by way of example.

5   **[0063]**   Let MPN\_L(j) and MPN\_H(k) represent pages in low and high memory Mem\_L, Mem\_H, respectively. For example, MPN\_L(j) might be any page in the lowest 4 GB of system memory, and MPN\_H(k) is any page with a hardware address above 4 GB. Using the transparent remapping method according to the invention, for each highly active page PPN(i), the mapping module 610, in particular, its remapping  
10   component 616, allocates a machine page in low memory (MPN\_L(j)), copies the data to be transferred from the original page in high memory (MPN\_H(k)) into the low page (MPN\_L(j)), and then changes the corresponding PPN→MPN mapping from PPN(i)→MPN\_H(k) to PPN(i)→MPN\_L(j). For the non-virtualized embodiment, or for the virtualized case in which the manager maintains a separate VPN→MPN map, the  
15   VPN→MPN mapping is changed from VPN(i)→MPN\_H(k) to VPN(i)→MPN\_L(j). The original page (MPN\_H(k)) can then be reclaimed using conventional mechanisms. Note that it is possible that a particular VPN or PPN might be already mapped to low memory for some other reason, that is, at least initially not as a result the remapping procedure of the invention; the invention may be used to manage mapping and remapping of such  
20   pages as well.

**[0064]**   In order for the manager to efficiently remap highly active, "hot" pages, it is necessary to determine just what these pages are. According to one method used in a prototype of the invention, the manager 605 includes, as part of the page activity module 614, a relatively small (in the prototype, 256 entries) hash table that is entered  
25   using the eight least significant (low-order) bits of PPNs used in I/O operations as the index. (Of course, a hash table of different size will be indexed using a different number of PPN bits.) The remaining higher order bits of each PPN thereby form its "tag." Each entry of the table includes the tag (or, alternatively, the full PPN) of the of the most recently issued PPN that hashes to the given index, the MPN to which the PPN is  
30   currently mapped (alternatively, a single bit that indicates whether it is currently mapped to high or low memory), and a count of the number of times this PPN has been issued.



The table is evaluated periodically and any PPN(i) whose count exceeds a predetermined threshold is then selected for remapping into low memory.

**[0065]** It is of course possible that two different PPNs will be issued that hash to the same entry in the table. Storing the tag of the most recently issued PPN allows the

5 page activity module 614 to avoid ambiguity, but it is still possible that a currently issued PPN will cause the entry of the earlier PPN to be overwritten. This method for measuring the activity of PPNs (using the count) therefore represents one choice in the trade-off between accuracy of activity measurement and ease of implementation -- the hash table is efficient and very easy to implement, but measures activity only

10 approximately, and only for at most the number of PPNs equal to the number of entries in the table. By maintaining counts only for pages in the table, this technique uses very

**[0066]** little space compared with techniques that store counts for all pages in the system, and will be acceptable in many actual implementations of the invention. The hash table-based technique for measuring PPN activity is not, however, the only choice  
15 in the trade-off of accuracy vs. simplicity.

**[0067]** Another method for measuring page activity according to the invention is for the page activity module 614 to compute, preferably for each PPN(i) in the map 612 (or for each VPN, if a separate list of VPNs is maintained), a corresponding activity score  $a(i)$  as a function of statistics of use of the respective PPN(i). In order to reduce space  
20 requirements, it would also be possible to calculate activity scores for only a subset of the pages in the system, for example, the subset of pages currently resident in a limited-size table such as a cache established for the purpose, or even a randomly selected subset of PPNs in use. The different values of  $a(i)$  may be maintained as a vector that augments the PPN→MPN map 612; any suitable, known indexing scheme may be used  
25 to associate  $a(i)$  with its respective PPN(i).

**[0068]** Activity may be measured in various ways, such as the number of I/O requests for page PPN(i) per unit time or over some time interval, or the number of times page PPN(i) must be copied, or the ratio of requests for I/O of PPN(i) relative to total I/O requests over some time interval, etc. Each of these measures of activity may  
30 be stored simply by augmenting the map 612 to include the score  $a(i)$  for each PPN(i). Note that time may be measured in either the real sense, that is, system time, or as a

function of CPU cycles, or as a function of the virtual time in which the VM is running, or according to any other predetermined definition.

**[0069]** Assume that the PPN→MPN mappings are such that all PPNs are initially mapped to high memory, that is each PPN(i) is initially mapped to some MPN\_H(k).

5 Whenever the activity score  $a(i)$  of page number PPN(i) exceeds a high-activity threshold value  $\tau_H$ , and assuming that not all low memory pages MPN\_L(j) are currently the target of a PPN→MPN\_L mapping, the remapping module 616 then changes the mapping of PPN(i) from high memory to low memory. In other words, if  $a(i) > \tau_H$ , then the initial mapping PPN(i)→MPN\_H(k) is changed to PPN(i)→MPN\_L(j), where  
10 MPN\_L(j) is any currently free MPN in low memory.

**[0070]** The value of the high-activity threshold  $\tau_H$  may be determined using conventional testing and design methods and will depend on such factors as the anticipated amount of I/O, the amount of low memory available, etc. Note that the threshold  $\tau_H$  may also be a parameter that is calculated based on the demand for low  
15 memory during some interval, on the proportion of low memory that is currently free, or on some combination of these or other factors. The threshold  $\tau_H$  for remapping into low memory may thus be made to rise/fall as the demand for low memory increases/decreases or as it becomes more used up/free.

**[0071]** The decision to remap a page into low memory Mem\_L of course increases  
20 the demand for low pages, which may be a scarce resource. The activity score  $a(i)$  for each PPN(i) is therefore preferably recalculated from time to time so that a PPN that has been remapped to low memory will not remain so mapped if the advantage of doing so decreases or disappears. A PPN→MPN\_L mapping should not be allowed to remain if the corresponding PPN is being used in I/O relatively infrequently, and low memory  
25 pages are scarce. Such recalculation may be performed periodically, after each or a certain number of I/O request(s) for the respective PPN(i) or for *any* mapped PPN, or according to some other predetermined schedule or according to other predetermined conditions. Note that a PPN(i) that is currently mapped into low memory should preferably be mapped back into high memory only if this would free up a low page that  
30 could be put to more productive use for another remapping, that is for an even more active PPN, in order to reduce aggregate copying. If plenty of low memory is available,

or there are no more deserving (that is, more active) remap candidates, then mapping back into high memory will not be necessary.

**[0072]** The level of scarcity of low memory may be measured in different ways. It will usually be impractical or even impossible, however, to determine how many pages of low memory are "active" and actually needed by a process. One way to deal with this is for the manager to always keep some amount of low pages in reserve and then to use a simple "low watermark" threshold to determine when to start reclaiming via low-to-high memory remapping.

**[0073]** Indirect and approximate estimates of scarcity may also be used. In order to determine one such indirect measure of low-memory availability the manager 605, the remapping or cost modules 616, 618 (for example) could keep track of how often requested low pages are unavailable. The ratio  $p$  of failed to total low-memory allocations can then be taken as an approximation of the degree of scarcity of low memory.

**[0074]** One way to avoid "overpopulation" of low memory is to reclaim low pages (via low-to-high memory remapping) randomly. According to this easily implemented procedure, when the number of available low memory pages drops below some threshold, then the remapping module 616 picks at random a PPN currently mapped to a low MPN, and then remaps it back to high memory in order to free up the low page.

Note that if the randomly selected PPN is actually being used heavily for I/O, then it will eventually be remapped back to low memory again.

**[0075]** Yet another procedure that avoids overpopulation of low memory is to remap all PPNs to high memory Mem\_H, reset all activity scores  $a(i)$  to zero (or whatever other initial value if assumed), and then to begin the activity-driven remapping to low memory Mem\_L anew. This may be done at set intervals, for example, or after some period of relative low I/O activity, after completion of some procedure, or at any other appropriate time. Although straightforward and "clean," this approach has the disadvantage of potentially remapping a large number of pages, some of which may still be involved in relatively frequent I/O requests.

[0076] Other low-to-high remapping approaches provide greater flexibility. For example, a PPN may be remapped back to high memory if its activity score falls below a low-activity threshold  $\tau_L$ , which may be determined using the same methods as  $\tau_H$ . It is preferable that  $\tau_L$  should be kept sufficiently less than  $\tau_H$  so that remapping is not done too often. This remapping "hysteresis" or "spread" between  $\tau_L$  and  $\tau_H$  may also be made a function of such factors as the current level of general I/O activity, the proportion of free MPN\_L(j), etc.

[0077] Still another optional mechanism for dynamic remapping from low back to high memory uses the scarcity ratio  $\rho$  defined above. According to this mechanism, either or both of the thresholds  $\tau_L$  and  $\tau_H$  are adjusted by a function of  $\rho$ . For example,  $\tau_H$  could be adjusted upward from some predetermined minimum value as a function of  $\rho$ ; similarly,  $\tau_L$  could be adjusted downward from some predetermined maximum value as another function of  $\rho$ . As more and more low memory becomes taken, this scheme would require an increasing level of activity in order for a particular PPN(i) to be mapped into the increasingly scarce low memory resource.

[0078] The activity scores  $a(i)$  may also be calculated in such a manner that they "decay" or grow more or less continuously; in other words, mappings into low memory may be allowed to "age." For example, let  $c(i,t)$  be a measure of the I/O activity of PPN(i) during the period  $(t, t-T)$ . Values for  $c(i,t)$  are then preferably also stored as a vector in the page activity module 614. This measure  $c(i,t)$  might be the number of times PPN(i) has been involved in I/O during the interval, some function of the percentage of total request I/O operations during the interval that involved PPN(i), etc. To allow for substantially continuous "aging" of mappings, after each interval  $T$ , the page activity module 614 could recalculate  $a(i)=a(i,t)$  as follows:

$$a(i,t+T) = \mu * a(i,t) + (1-\mu) * c(i,t+T) \quad (\text{EQ 1})$$

where  $\mu$  is a constant that may be chosen using known theoretical and experimental methods to provide a desired degree of "memory" and, preferably, to keep  $a(i)$  always within some normalized range. The more frequently PPN(i) is involved in I/O operations, the larger its activity score will grow to be. Its score will drop, however, as the frequency of use of PPN(i) decreases. Using such an aging function also has the

effect of smoothing out the frequency of remapping operations, but it therefore also introduces a lag that will depend on  $\mu$ .

**[0079]** Other aging functions may of course also be used. For example, whereas the formula for calculating  $a(i,t)$  shown above is recursive and takes into account all I/O activity for PPN(i) that has taken place since the scores  $a(i)$  were most recently reset, it would also be possible to use an aging function of only the current and the previous  $m$  activity intervals. Thus, one could also calculate the activity score for a particular PPN(i) as follows:

$$a(i,t) = \mu_0 * c(i,t) + \mu_1 * c(i,t-T) + \mu_2 * c(i,t-2T) + \dots + \mu_m * c(i,t-mT) \quad (\text{EQ 2})$$

where  $m$  and the weights  $\mu_0, \mu_1, \mu_2, \dots, \mu_m$  may be chosen using normal design methods. The advantage of the recursive formula in EQ 1, however, is that it requires only two storage spaces per PPN(i) and fewer additions and multiplications than the  $(m+1)$ -element formula EQ 2.

**[0080]** In addition to (or instead of) remapping pages from high to low memory (or vice versa) as a function of activity, remapping may also be made contingent on cost: According to this optional aspect of the invention, for a given page PPN(i), the manager, for example, a cost-evaluation module 618 (Figure 2), keeps track of the amount of copying overhead incurred while doing I/O for that page, and then remaps when this copying overhead exceeds the actual remap cost (or some function of the actual remap cost). Note that it will generally be easy to estimate the cost of a page copy, since it is nearly constant in most cases.

**[0081]** The cost-evaluation module 618 may, additionally, also be used to implement a "global" decision mechanism that determines whether any remapping should be performed or continued at all. According to this aspect of the invention, the module 618 evaluates the current cost of remapping and if this cost rises above a predetermined threshold, then page remapping is at least temporarily discontinued until the cost function once again drops below the threshold. While remapping is suspended, I/O requests to "high" memory may be carried out by copying the data for each requested page (for *write* operations), each time it is requested, to the low-memory buffer 620 as

in the prior art. (For *read* operations it is necessary only to allocate a buffer page, possibly with partial page copying as mentioned above, and to copy the input data to the subsystem that requested the read.)

**[0082]** As is mentioned above, in systems such as those with the Intel x86

5 architecture, a page is the smallest unit of memory than can typically be remapped. I/O requests, however, may be of any size, that is, less (even much less) than a page, exactly a page, or larger (even much larger) than a page (which also means it will occupy more than one page). For example, most DMA uses what are known as "scatter-gather lists" for an I/O request, where each element of a list is an <address,  
10 length> pair.

**[0083]** In cases where an I/O request spans more than one page, the corresponding pages may be (but need not be) contiguous (that is, numerically adjacent) as PPNs; however, depending on the implementation, the MPNs these PPNs map to may not be contiguous. The remapping techniques of the invention described above may in such  
15 cases also be used to ensure mapping to contiguous MPNs. For example, a single element for an 8K write to contiguous PPNs might be converted to two 4K writes to their corresponding noncontiguous MPNs. The manager 605 could then use the remapping procedure according to the invention to copy/remap so that the element is contiguous in MPN space as well in PPN space. In short, the invention may be used to remap non-  
20 contiguous pages to contiguous pages. This would then avoid the need for page splits on future transfers and will improve performance accordingly.

**[0084]** Although page remapping as such is a known technique with a variety of applications, those skilled in the art of operating system design will now appreciate that several aspects of its implementation according to the invention are not only novel but  
25 also provide particular advantages. First, according to the invention, page remapping is used in a manner that helps avoid expensive copy operations while performing I/O operations to devices that have limited addressing capabilities. Second, in the preferred, virtualized embodiment of the invention shown in Figure 3, the invention exploits the extra level of indirection made possible by virtualization in order to  
30 transparently remap (virtualized) "physical" memory in the guest OS that is running in the VM. By way of contrast, note that in a conventional OS, it is not possible to remap

"physical" memory at all, which is accessed directly by devices via DMA. Finally, the invention provides for an *adaptive* remapping process that uses dynamic statistics to reduce the need for copying of pages.